

**This Page Is Inserted by IFW Operations
and is not a part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- **BLACK BORDERS**
- **TEXT CUT OFF AT TOP, BOTTOM OR SIDES**
- **FADED TEXT**
- **ILLEGIBLE TEXT**
- **SKEWED/SLANTED IMAGES**
- **COLORED PHOTOS**
- **BLACK OR VERY BLACK AND WHITE DARK PHOTOS**
- **GRAY SCALE DOCUMENTS**

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**



Europäisches Patentamt
European Patent Office
Office européen des brevets



Publication number: 0 620 522 A2

EUROPEAN PATENT APPLICATION

Application number: 94301785.5

Int. Cl.⁵: G06F 9/445

Date of filing: 14.03.94

Priority: 13.04.93 US 46827

Date of publication of application:
19.10.94 Bulletin 94/42

Designated Contracting States:
DE FR IT NL SE

Applicant: SUN MICROSYSTEMS, INC.
2550 Garcia Avenue
Mountain View, California 94043-1100 (US)

Inventor: Nelson, Michael N.
158 Hillcrest Road
San Carlos, California 94070 (US)
Inventor: Hamilton, Graham
3143 David Court
Palo Alto, California 94303 (US)

Representative: Wombwell, Francis
Potts, Kerr & Co.
15, Hamilton Square
Birkenhead Merseyside L41 6BR (GB)

High performance dynamic linking through caching.

A method is disclosed for a method to dynamically link a new program image and related library programs into an executable application program image. The method provides for producing a linked list of the required programs by giving the linker a n argument representing the designated program image and a naming context which contains data on the associated library programs which are to be linked together. The linker finds all of the required programs, and links them together. The parent maps the program images into the designated addresses thereby completing the linking of the executable application program. In finding the required programs, the linker first checks the image cache to see if the new program and its related library programs are already linked and cached because it was executed before. If the new program is not found in the image cache, the linker object then checks the library program cache to see if the library programs are cached in partially linked form, and will use as many of these as it can find. For any library programs that must still be located, the linker will retrieve them from a data store and will then proceed to link all library programs and the new program image together to form an executable whole. This process of caching new program images with their library programs and caching partially linked library programs individually guarantees a procedure which minimizes the time delay in program start-up when a new program is executed.




FIG. 7

EP 0 620 522 A2

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to the fields of distributed computing systems, client-server computing and object oriented programming.

BACKGROUND

A computer programmer writes a program, the source program, in a high level language which typically makes use of symbols for addresses, and special characters or acronyms for operation codes. This source code comprises a set of instructions for the computer and data with which or upon which the instructions are to operate. These instructions and data must be loaded into a computer's memory at certain addresses in order for the computer to execute the program process. In order to make this happen, the source code is processed by a compiler which generates binary object code which a computer can execute. Before the computer can execute this newly written program, the program must go through several additional steps, during which the addresses in the program may be represented in different ways. The compiler typically will bind the symbolic addresses of the source code to relocatable addresses (such as 16 bytes from the beginning address of the program). A *linkage editor or loader* will in turn bind these relocatable addresses to absolute addresses (such as memory location 64216). Each binding is a mapping from one address space to another. This *binding* of instructions and data to memory addresses can be done either at compilation time, load time or at program execution time.

Referring to **Figure 1**, the several steps described above are depicted. The source program 1 is processed by a compiler 2 producing an object code module 3. This object code module 3, generally along with other previously compiled object modules 4, is processed by a linkage editor 5 to produce a load module 6. The load module 6 and any system libraries 7 required are processed by a loader 8 producing an in-memory binary image 10 of the original program and its related modules and libraries. This in-memory binary image 10 can now be executed by the computer.

Continuing to refer to **Figure 1**, when it is known at compile time 12 where the program will reside in memory, these programs are compiled with absolute code for addresses. In most cases however, it is not known at compile time where the program will reside in memory, and the compiler must generate relocatable code for the memory addresses. In this case, final binding of memory locations to the addresses is delayed until either

load time 13 or execute time 14. Some modern systems delay such address binding until execute time 14 when the program image can be moved during its execution from one memory segment to another or where the program start-up cost is not excessive because the program images contain few relocatable addresses, such as with position independent code (PIC). PIC is code generated by some compilers which can be placed anywhere in memory because all memory references are made relative to the *program counter*.

Modern computer operating systems are designed to optimize the use of memory space and to minimize user wait time. This is done in the address binding/program loading process, by *dynamic loading* of program object modules only when they are actually called by another module and *dynamic* linking of an object module to its system library routines only when they are required. In these cases the main program is loaded into memory and executed and supporting object modules or system libraries are not loaded unless they are called by the main program, thus saving memory space and load time at the expense of some program start-up time. Also, system libraries which will likely already be resident in memory can be dynamically linked to the executing main program when called, thereby not requiring a copy of the system libraries to be linked and loaded with each main program at load time, again saving memory space but at the cost of some program start-up time and some execution time. With dynamic linking, a *stub* is included in the image for each library-routine reference. This stub is a small piece of code that indicates how to locate the appropriate memory-resident library routine. When the stub is executed it replaces itself with the address of the routine and executes the routine. Under this scheme all programs that use a library routine use the same copy of the library code.

In order to take maximum advantage of dynamic linking and loading, program compilers must be designed to produce the necessary relocatable address references, the sub-routine stub code, and efficient PIC code. Unfortunately, existing compilers for object-oriented program modules cannot generate such code efficiently. For example, the Cfront 3.0 preprocessor and the G++ compiler generate virtual function tables as initialized data structures which are full of references to relocatable symbols. Thus the number of relocatable symbols in object-oriented program modules is much higher than in more traditional program modules and the program startup delay required to dynamically link these modules can rise to unacceptable levels due to the number of relocations. Therefore what is required is a system that provides efficient dynamic linking of program modules with large numbers of re-

locatable symbols.

The present invention provides an elegant solution to this problem by caching linked program images and also caching partially linked library programs.

SUMMARY OF THE INVENTION

The present invention fills this need for minimizing the system delay caused by linking new object oriented programs by establishing an elegant and efficient system of caching fully bound program images along with their system library modules, and establishing a second level of caches for caching relocatable system library programs whose addresses are provisionally fixed-up (partly linked), and using these caches to supply fully linked or partially linked modules when a call to another program/routine is made. Thus if a new application is to be linked, the system first checks the image cache to see if the program has been linked before and if so the linked version of the program along with its linked library programs is used without further processing overhead. If the new application is not found in the image cache, then an attempt is made to minimize the overhead by seeing if some or all of the related library programs have already been provisionally fixed-up (partly linked), by checking a library program cache. Only if a library program is not found in the library program cache, does the linker system need to do a complete program load and link operation, thus minimizing the linking overhead.

A method is disclosed for a method to dynamically link a new program image and related library programs into an executable application program image. The method provides for producing a linked list of the required programs by giving the linker an argument representing the designated program image and a naming context which contains data on the associated library programs which are to be linked together. The linker finds all of the required programs, and links them together. The parent maps the program images into the designated addresses thereby completing the linking of the executable application program. In finding the required programs, the linker first checks the image cache to see if the new program and its related library programs are already linked and cached because it was executed before. If the new program is not found in the image cache, the object then checks the library program cache to see if the library programs are cached in partially linked form, and will use as many of these as it can find. For any library programs that must still be located, the linker will retrieve them from a data store and will then proceed to link all library programs and the new program image together to form an execut-

able whole. This process of caching new program images with their library programs and caching partially linked library programs individually guarantees a procedure which minimizes the time delay in program start-up when a new program is executed.

DESCRIPTION OF THE DRAWINGS

The objects, features and advantages of the system of the present invention will be apparent from the following description in which:

Figure 1 illustrates the typical steps in compiling and loading a new program (Prior Art).

Figure 2 illustrates the SPRING operating system concept of an Object.

Figure 3 illustrates a High level flow chart of the invention.

Figure 4 illustrates a detailed flow chart of the invention.

Figure 5 illustrates a non-fixed-up program.

Figure 6 illustrates a fixed-up version of the program shown in **Figure 5**.

Figure 7 illustrates an example of a SPRING dynamic linking environment.

NOTATIONS AND NOMENCLATURE

The detailed descriptions which follow may be presented in terms of program procedures executed on a computer or network of computers. These procedural descriptions and representations are the means used by those skilled in the art to most effectively convey the substance of their work to others skilled in the art.

A procedure is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. These steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It proves convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like. It should be noted, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities.

Further, the manipulations performed are often referred to in terms, such as adding or comparing, which are commonly associated with mental operations performed by a human operator. No such capability of a human operator is necessary, or desirable in most cases, in any of the operations described herein in which form part of the present invention; the operations are machine operations.

Useful machines for performing the operations of the present invention include general purpose digital computers or similar devices.

The present invention also relates to apparatus for performing these operations. This apparatus may be specially constructed for the required purposes or it may comprise a general purpose computer as selectively activated or reconfigured by a computer program stored in the computer. The procedures presented herein are not inherently related to a particular computer or other apparatus. Various general purpose machines may be used with programs written in accordance with the teachings herein, or it may prove more convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these machines will appear from the description given.

DESCRIPTION OF THE PREFERRED EMBODIMENT

In the following descriptions for purposes of explanation, specific data and configurations are set forth in order to provide a thorough understanding of the present invention. The preferred embodiment described herein is implemented as a portion of the SPRING Object-Oriented Operating System created by Sun Microsystems®, Inc. (Sun Microsystems is a registered trademark of Sun Microsystems, Inc.) However, it will be apparent to one skilled in the art that the present invention may be practiced without the specific details and may be implemented in various computer systems and in various configurations, or makes or models of tightly-coupled processors or in various configurations of loosely-coupled multiprocessor systems. Moreover, it will be clear to those skilled in these arts that the present invention may be implemented in a non-object oriented computing system.

A SPRING *domain* is an *address space* with a collection of *threads*. A given domain may act as the server of some objects and the clients of other objects. The implementor or object manager and the client can be in the same domain or in a different domain.

The spring object model

SPRING has a slightly different way of viewing objects from other distributed object oriented systems and it is necessary to clarify this before discussing the details of the present invention.

Most distributed systems present a model wherein objects reside at server machines and client machines possess object handles that point to the object at the server. (See figure 2a.) Clients pass around object handles rather than objects.

SPRING presents a model wherein clients are operating directly on objects, not on object handles. (See figure 2b.) Some of these objects happen to keep all their internal state at some remote site, so that their local state merely consists of a handle to this remote state. An object can only exist in one place at a time, so if an object is transmitted to someone else then the transmitter of the object ceases to have the object. However, the object can be copied before being transmitted, which might be implemented such that there are now two distinct objects pointing to the same remote state.

So whereas in systems such as MACH, one might talk of several clients having object handles that reference some remote object, in SPRING one would talk about several clients having objects that reference the same remote state.

Dynamic Linking of Process Images In SPRING

Referring now to Figure 3 the basic steps of the present invention are briefly described. When a client domain in SPRING wishes to start a program, the client domain executes a "Link (Image Memory Object, Naming Context Object)" command on a linker object. The linker object first checks his image cache 22 by looking-up the image memory object to see if it contains a fully-linked version of the program represented by the image memory object and that this cached version was linked using the library programs in the naming context object. If a match is found 28 the linker object is basically finished with its work and merely returns the cached list of memory objects and related addresses 38 which are associated with the found image memory object. If the check of the image cache produces no match 24 then the linker object finds a list of names of the library programs required by the image memory object and using the naming context object, obtains a memory object associated with each library program, and proceeds to check the library program cache for a match for each library program 26. If a match is found for a library program memory object 36 the linker object obtains the memory object representing the provisionally fixed-up copy of the found library program and the associated address at which this program is to be mapped. If no match is found 30, the linker object decides on an address at which to map the library program in question, creates a memory object that is a copy-on-write copy of the library program in question and maps this memory object at the address it selected for it. The linker object further provisionally fixes-up this new copy of the program and stores the library program's memory object and related address in the library program cache 32. After finding a mem-

ory object for each library program required, whether they were found in the library program cache or new copies created, the linker object now resolves all unresolved symbol references in the new program image and in all of these library programs and writes a copy of this fully-linked version into the image cache 34. Lastly the linker object returns to the client domain a list of the memory objects and related addresses for this fully-linked new program image and its associated library programs.

The above description is a summarized description of the present invention which is explained in more detail in the flow chart depicted in Figure 4 and described in examples of the process below. Those skilled in the art will recognize that the use of names such as linker object, client domain etc. are for illustrative purposes and that the invention may be practiced by programs which may be called any kind of name.

Before proceeding to explain the invention in more detail, it is useful to describe an example of a "non-fixed-up program" and what happens as the relocations occur to make it a "fixed-up" or executable program image. Referring now to Figure 5 a non-fixed-up program is depicted. File F 40 contains a program image. The file has four parts: a header 42 that contains the address 44 to begin executing the program and the list of shared library programs 46 to use; the code for the program 52; a symbol table 48 that contains all external symbols defined in the program; and a relocation table 50 that contains a list of what symbols to relocate. In this example, the program header 42 says that the program needs shared library SL 46. The relocation table 50 shows that the program code 52 contains an instruction at address A1 that references an external symbol *foo* 54. In addition, the symbol table 48 shows that the program defines the external function *bar* 56 at address A2. Also shown in Figure 5 is file F2 containing shared library SL 60. File F2 has the same four parts as File F: a program header 62; the code for the library program 72; a symbol table 68; and a relocation table 70. Since this shared library does not depend on any other shared libraries, the shared library list is empty 66. The relocation table 70 shows that the library code contains an instruction at address B1 that references the external symbol *bar* 64. In addition the symbol table 68 shows that the library code defines the external function *foo* at address B2 76. Figure 6 shows a "fixed-up" version of the same program and library program which is depicted in Figure 5. Figure 6 shows an address space 80 containing the fixed-up program code 82 and the fixed-up library code 84. The code for the program image in File F (52 in Figure 5) is mapped at the address 86 that is fixed for all

program images. The value of this address is such that each code address in the image (52 in Figure 5) is correct (for example, the code at address A1 (58 in Figure 5) is mapped at address A1 in the address space 86). The shared library SL on the other hand is mapped beginning at address space location address R1 and thus each library code address is offset by R1 90 in the address space 80. The external reference to the symbol *foo* in the instruction A1, which was shown in the File F relocation table (54 in Figure 5) to be at instruction A1, was satisfied by the shared library. The symbol *foo* was defined at address B2 in the shared library (76 in Figure 5) and since address B2 was mapped to address R1 + B2 in the Address space 80, the actual address put in A1 for *foo* was R1 + B2 88. The external reference to the symbol *bar* in the instruction B1 shown by the library relocation table (64 in Figure 5) was satisfied by the program code which defined *bar* at address A2 in the program image (56 in Figure 5). The instruction at address B1 in the shared library is actually mapped to address R1 + B1 in the address space 90 and so the actual address placed in the instruction at R1 + B1 is A2 92.

Referring now to Figure 7 an environment in an exemplary SPRING object oriented system is depicted with which to illustrate the preferred embodiment of the invention. Shown are: a Naming Service 100 containing an implementor for context object C 102; a Domain D 104 containing an Address Space 106; a Client Domain 108 containing a mapping list 120 and objects C 109, L 110, M 112, AS 114, D 116, and DM 118; a SPRING kernel 122, containing a Domain Manager 124 which itself contains an implementor of domain object D 126 which is shown containing object AS 128; a Virtual Memory Manager 130 also contained in the SPRING kernel which contains an address space implementor for objects AS 132; a File Server Domain 152 containing implementors for memory object M 154, memory object M2 156, memory object SL 158, memory object SL2 160 and memory object SL3 162; and finally a Linker Domain 140 containing an implementor for linker object L 142, an image cache 144 containing an exemplary cache entry 146 which itself contains a list of memory objects and related addresses 172 and Memory objects C 164, M 166, M2 168 and SL3 170, and a library cache 148 containing an exemplary cache entry 150 which itself contains memory objects SL 172 and SL2 174, and the cached program starting address 176.

It should be noted that the following objects are equivalent: M 112 and M 166; M2 175 and M2 168; SL3 177 and SL3 170; and C 109 and C 164. Two objects for the purpose of this invention are considered equivalent if they reference the source stat

on the server.

In the examples which follow, a client domain wishes to start a program named *foo*. The configuration after the program is started is shown in **Figure 7**.

The simplest example is where there is a cache hit in the image cache. In this case starting the new domain goes through the following steps:

- o The client domain 108 looks up *foo* via the Spring naming service 100 and gets back a memory object M 112 that stores the program image for *foo*. The program image stored in M needs shared library SL in order to be dynamically linked; the names of the libraries that an image needs to be linked are stored inside memory object M 112 whose implementation is on disk 154.
- o The client domain 108 looks up a linker object L 110 implemented by a linker domain 140 using the naming service 100.
- o The client domain 108 invokes the link method on L 110 passing in two parameters: memory object M 112 and a context object C 109 implemented by the naming service 100.
- o The linker 140 looks in its cache using M 112 and C 109 as its key and discovers that it has an entry with key M 166 and C 164 that matches (M 166 is equivalent to M 112 and C 164 is equivalent to C 109). The cache entry contains a cached copy of a fully-linked version of M 166 that was linked using C 164. The fully-linked version consists of two memory objects: M2 168 and SL3 170. M2 168 is a fully linked copy of M 166 linked at address A1 (the address to map program images is fixed for all programs to be the value A1) and SL3 170 is a fully linked copy of the memory object for the shared library SL 172 linked at address A2 (the list of memory objects and corresponding linked addresses is shown at 173).
- o The linker domain 140 returns a list of (address to map, memory object pairs) 173 to the client domain 108 (this list is labeled mapping list 120 in the client domain 108). The first entry in the list 121 contains address A1 and memory object M2. The second entry in the list 123 contains address A2, and memory object SL3.
- o The client domain 108 looks up a domain manager object DM 118 using the naming service 100.
- o The client domain 108 invokes the *create_domain* method on object DM 118 and gets in return a new domain object D 116.
- o The client domain 108 invokes the *get_address_space* method on D 116 and gets an object AS 114 that represents D's address

space 132.

- o The client domain 108 invokes the *copy_and_map* method on the AS object 114 to map object M2 175 at address A1 and SL3 177 at address A2 copy-on-write. (Note that this information was obtained from the mapping list 120).
- o The client domain 108 starts domain D 126 running by invoking the *enter* method on domain object D 116.

Note that in **Figure 7**, the library cache 148 also has a cached copy of shared library SL 172 and a provisionally fixed-up version which is in memory object SL2 174. However, since we got a hit on the image cache 144 we didn't need to consult the library cache 148.

A second example is when there is no image cache hit but there is a library cache hit. For this example assume that the image cache 144 in **Figure 7** is empty when the client domain 108 tries to link the program image. The linker domain 140 does the link using the following steps:

- o The linker domain 140 looks up M 112 and C 109 in its image cache 144 and does not find a match.
- o The linker domain 140 looks in memory object M 112 and discovers that it needs a shared library program named SL to be linked.
- o The linker domain 140 looks up the name SL using context object C 109. (Note that this step would be repeated for each library program required). The result is the memory object SL 300 (this memory object is not shown in the linker domain 140) that contains the contents of the shared library program SL.
- o The linker domain 140 looks in its library cache 148 using memory object SL 300 as an argument, and discovers that it has an entry with key SL 172 that watches (i.e., SL 172 is equivalent to SL 300). The cached entry contains a cache copy of a provisionally fixed-up version of shared library SL which is memory object SL2 174.
- o The linker domain 140 creates a memory object M2 168 that is a copy-on-write copy of M 112 and creates a memory object SL3 170 that is a copy-on-write copy of SL2 174.
- o The linker domain 140 resolves all unresolved references in M2 168 using SL3 170 and resolves all unresolved references in SL3 170 using M2 168 (how this is done was described above and is depicted in the flow chart). The result is that M2 168 and SL3 170 comprise a fully-linked version of program *foo*. Again, it should be noted that if there were multiple library programs required, the

same steps would be performed for each one.

- o The linker domain enters M2 168 and SL3 170 into its image cache 144 along with their related keys (memory object C 164 and memory object M 166) and returns a list of (address to map, memory object pairs) 173 to the client domain 108 (this list is labeled mapping list 120 in 108). The first entry in the list 121 contains address A1 and memory object M2 and the second entry in the list 123 contains address A2 and memory object SL3.

The client domain 108 then follows the same steps in the first example to create the new domain.

The last example is when there is no image or library cache hit. For this example assume that when the client domain 108 tries to link the program image, the image cache 144 and the library cache 148 are empty. The linker domain 140 does the link using the following steps:

- o The linker domain 140 looks up M 112 and C 109 in its image cache 144 and does not find a match.
- o The linker domain 140 finds memory object M 112 and maps it into its address space, and discovers that M 112 needs a shared library named SL to be linked.
- o The linker domain 140 looks up the name SL using context object C 109. The result is the memory object SL 172 that contains the contents of the shared library program SL.
- o The linker domain 140 looks in its library cache 148 using memory object SL 172 as the argument, and does not find a match (recall that in this example, we assumed the library cache to be empty).
- o The linker domain 140 decides that memory object SL 172 must be fixed-up at an available address in its address space (call it A2).
- o The linker domain 140 creates a memory object SL2 160 that is a copy-on-write copy of SL 172 and maps SL2 into its address space.
- o The linker domain 140 provisionally fixes-up SL2 (see the discussion on fix-ups above and the detailed flow-chart in Figure 4 for details) and stores SL2 174 in the library cache 148 - (along with its key SL 172).
- o The linker domain 140 creates a memory object M2 168 that is a copy-on-write copy of M 112 and creates a memory object SL3 170 that is a copy-on-write copy of SL2 174.
- o The linker domain 140 resolves all unresolved references in M2 168 using SL3 170 and resolves all unresolved references in SL3 170 using M2 168 (how this is done is explained

in the discussion on fix-ups above and in the detailed flow chart in Figure 4). The result is that M2 168 and SL3 170 comprise a fully-linked version of program foo.

- o The linker domain 140 enters M2 168 and SL3 170 into its image cache 144 (along with its key, memory object M166 and its context object C 164) and returns a list 172 of (address to map, memory object pairs) to the client domain 108. The first entry of the list 121 contains address A1 and memory object M2. The second entry in the list 123 contains address A2 and memory object SL3.

The client domain 108 then follows the same steps as above in the first example, to create the new domain 104.

An alternative embodiment of the invention is to use only the shared library program cache in the dynamic linker system, following the same steps as outlined above starting with the second example.

While the above examples describe the presently preferred embodiment, and describe specific program domains and a certain sequence of steps, those skilled in these arts will recognize that these specifics are not essential to the practice of the invention which basically embodies the use of caches for retaining fully linked images of programs and their associated libraries as well as partially fixed-up library routines with which to minimize the overall dynamic linking overhead cost.

Claims

1. In a computing system wherein there exists an application program, library programs and an operating system, a method to dynamically link a designated program image and related library programs into an executable application program image, said method comprising the steps of:

invoking a link method with parameters representing said designated program image and parameters representing a naming context for said related library programs;

locating said designated program image and said related library programs;

linking said related library programs with each other and with said designated program image, and creating a list of program identifiers and related addresses comprising an identifier and related address for said designated program image and for each of said related library programs; and

mapping said designated program image and each of said related library programs at said related addresses from said list of program identifiers,

whereby said designated program image

and said related library programs are dynamically linked into an executable program image.

2. The method of claim 1 wherein the step of locating said designated program image and said related library programs comprises a step of selecting said designated program image and said related library programs from an image cache in the event said designated program image is found in said image cache.
3. The method of claim 2 wherein the step of locating said designated program image and said related library programs further comprises the steps of making a copy of said designated program image in the event said designated program image is not found in said image cache, and selecting said library programs from a library program cache in the event said library programs are found in said library program cache.
4. The method of claim 3 wherein the step of locating said designated program image and said related library programs further comprises the following steps:
 - selecting an address to which to link a required library program in the event said library programs are not found in said library program cache;
 - making a copy of said required library program;
 - provisionally fixing-up relocation symbols in said copy of said required library program; and
 - adding said copy of said required library program to said library program cache.
5. The method of claim 4 wherein said step of linking said related library programs with each other and with said designated program image, and creating a list of program identifiers and related addresses comprising an identifier and related address for said designated program image and for each of said related library programs further comprises the steps of:
 - adding to said image cache a copy of said designated program image with a copy of said designated program's related library programs and said list of program identifiers and related addresses for each of said programs.
6. The method of claim 5 wherein all programs are object oriented programs.
7. The method of claim 1 wherein the step of locating said designated program image and said related library programs further comprises

a step of selecting said library programs from a library program cache in the event said library programs are not found in said library program cache.

8. The method of claim 7 wherein the step of locating said designated program image and said related library programs further comprises the following steps:
 - selecting an address to which to link a required library program in the event said library programs are not found in said library program cache;
 - making a copy of said required library program;
 - provisionally fixing-up relocation symbols in said copy of said required library program; and
 - adding said copy of said required library program to said library program cache.
9. The method of claim 8 wherein said step of linking said related library programs with each other and with said designated program image, and creating a list of program identifiers and related addresses comprising an identifier and related address for said designated program image and for each of said related library programs further comprises the step of:
 - adding to said image cache a copy of said designated program image with a copy of said designated program's related library programs and said list of program identifiers and related addresses for each of said programs.
10. The method of claim 9 wherein all programs are object oriented programs.
11. In a computing system, a method for dynamically linking a designated program image and related library programs into an executable program image, said method comprising the steps of:
 - invoking a link method with parameters representing said designated program image and parameters representing library programs related to said designated program image;
 - checking an image cache and in the event that a key is found in said image cache which matches said designated program image and said related library programs, obtaining a list of program identifiers and related addresses from said image cache; and
 - mapping each program identified by said program identifiers at said related address from said list of program identifiers, whereby said designated program image and said related library programs are dynam-

ically linked into an executable program image.

12. The method of claim 11 comprising the additional steps of:

creating a first list containing names of said library programs which are related to said designated program image in the event that an entry is not found in said image cache which matches said designated program image and said related library programs; and

resolving each of said names of said library programs contained on said first list into a library program identifier and adding each of said library program identifiers to a second list.

13. The method of claim 12 comprising the additional steps of:

checking a library program cache for a match for each library program identifier in said second list;

obtaining a matching library program from said library program cache in the event that a matching library program identifier is found; and

adding said matching library program identifier from said library program cache to a third list.

14. The method of claim 13 comprising the additional steps of:

selecting an address to which to link a required library program identified in said second list in the event that a matching library program identifier is not found in said library program cache;

making a copy of said required library program;

provisionally fixing-up relocation symbols in said copy of said required library program;

adding said provisionally fixed-up copy of said required library program to said library program cache; and

adding said provisionally fixed-up library program identifier to said third list.

15. The method of claim 14 comprising the additional steps of:

making a copy of said designated program image and adding an identifier of said copy of said designated program image to a fourth list, said copy of said designated program image having a symbol table and a relocation table; and

making a copy of each of said library programs from said third list and adding an identifier of said copy of said library programs to said fourth list, each of said copies of said library programs having a symbol table and a

relocation table.

16. The method of claim 15 comprising the additional steps of:

finding each symbol to be relocated in said copy of said designated program image; and

searching all of said library programs identified in said fourth list and in the event that a first symbol definition which matches said symbol to be relocated is found, using an address of said first symbol definition which matches said symbol to be relocated, to perform a relocation for said symbol to be relocated.

17. The method of claim 16 comprising the additional steps of:

using said symbol table in each of said library programs identified in said fourth list in turn, starting with a first of said library programs identified in said fourth list, comparing each symbol of said symbol table to said designated program image's symbol table and marking a symbol in said library program's symbol table as overridden in the event that said symbol matches a symbol in said designated program image's symbol table;

comparing each symbol of each of said symbol tables to all of said symbol tables of said library programs identified in said fourth list which are higher in said fourth list than said symbol table being checked and which have been marked as already checked;

upon completion of the check of said symbol table of each of said library program identified in said fourth list, marking said library program's identifier in said fourth list as one whose symbol table has been checked; and

comparing said symbols in said symbol tables in each of said library programs identified in said fourth list in like manner until all of said symbol tables have been compared.

18. The method of claim 17 comprising the additional steps of:

using said relocation table in each of said programs identified in said fourth list in turn, starting with a first of said programs identified in said fourth list, scanning each of said relocation tables to find an undefined symbol;

in the event that a symbol in one of said relocation tables is a defined symbol, checking said defined symbol and in the event said defined symbol has been marked as overridden, undoing a relocation address for said defined symbol which has been marked as overridden;

using each undefined symbol and each defined symbol whose relocation address was undone, as a search argument, checking said symbol table of said copy of said designated program image and said symbol tables of each of said library programs identified in said fourth list to find a symbol matching said search argument, and in the event said matching symbol is found, using said found symbol's address to perform said address relocation for said symbol which was used as said search argument; and

raising an exception message in the event no symbol is found to match said search argument.

19. The method of claim 18 comprising the additional steps of:

creating a cachable list of program identifiers, containing a copy of said designated program image and each of said library programs identified in said fourth list;

entering said cachable list of program identifiers into said image cache; and

returning said cachable list of program identifiers in response to said link method invocation.

20. The method of claim 19 wherein said programs are objects and said computing system is an object oriented system.

21. The method of claim 20 wherein said identifiers and said program identifiers are memory objects.

22. In a computing system, a method for dynamically linking a designated program image and related library programs into an executable program image, said method comprising the steps of:

invoking a link method with parameters representing said designated program image and parameters representing library programs related to said designated program image;

checking a library program cache and in the event that said library program cache contains said related library programs;

creating a list of library program identifiers and related addresses from said library program cache;

linking each of said library programs to each other and to said designated program image; and

mapping each program identified by said program identifiers at said related address from said list of program identifiers,

whereby said designated program image

and said related library programs are dynamically linked into an executable program image.

23. The method of claim 22 wherein said step of creating a list of library program identifiers and related addresses from said library program cache comprises the additional steps of:

creating a first list containing names of said library programs which are related to said designated program image; and

resolving each of said names of said library programs contained on said first list into a library program identifier and adding each of said library program identifiers to a second list.

checking said library program cache for a match for each library program identifier in said second list;

obtaining a matching library program from said library program cache in the event that a matching library program identifier is found; and

adding said matching library program identifier from said library program cache to a third list.

24. The method of claim 23 comprising the additional steps of:

selecting an address to which to link a required library program identified in said second list in the event that a matching library program identifier is not found in said library program cache;

making a copy of said required library program;

provisionally fixing-up relocation symbols in said copy of said required library program;

adding said provisionally fixed-up copy of said required library program to said library program cache; and

adding said provisionally feed-up library program identifier to said third list.

25. The method of claim 24 comprising the additional steps of:

making a copy of said designated program image and adding an identifier of said copy of said designated program image to a fourth list, said copy of said designated program image having a symbol table and a relocation table; and

making a copy of each of said library programs from said third list and adding an identifier of said copy of said library programs to said fourth list, each of said copies of said library programs having a symbol table and a relocation table.

26. The method of claim 25 comprising the additional steps of:

finding each symbol to be relocated in said copy of said designated program image; and

searching all of said library programs identified in said fourth list and in the event that a first symbol definition which matches said symbol to be relocated is found, using an address of said first symbol definition which matches said symbol to be relocated, to perform a relocation for said symbol to be relocated.

27. The method of claim 26 comprising the additional steps of:

using said symbol table in each of said library programs identified in said fourth list in turn, starting with a first of said library programs identified in said fourth list, comparing each symbol of said symbol table to said designated program image's symbol table and marking a symbol in said library program's symbol table as overridden in the event that said symbol matches a symbol in said designated program image's symbol table;

comparing each symbol of each of said symbol tables to all of said symbol tables of said library programs identified in said fourth list which are higher in said fourth list than said symbol table being checked and which have been marked as already checked;

upon completion of the check of said symbol table of each of said library program identified in said fourth list, marking said library program's identifier in said fourth list as one whose symbol table has been checked; and

comparing said symbols in said symbol tables in each of said library programs identified in said fourth list in like manner until all of said symbol tables have been compared.

28. The method of claim 27 comprising the additional steps of:

using said relocation table in each of said programs identified in said fourth list in turn, starting with a first of said programs identified in said fourth list, scanning each of said relocation tables to find an undefined symbol;

in the event that a symbol in one of said relocation tables is a defined symbol, checking said defined symbol a second time and in the event said defined symbol has been marked as overridden, undoing a relocation address for said defined symbol which has been marked as overridden;

using each undefined symbol and each defined symbol whose relocation address was

undone as a search argument, checking said symbol table of said copy of said designated program image and said symbol tables of each of said library programs identified in said fourth list to find a symbol matching said search argument, and in the event said matching symbol is found, using said found symbol's address to perform said address relocation for said symbol which was used as said search argument; and

raising an exception message in the event no symbol is found to match said search argument.

29. The method of claim 28 comprising the additional steps of:

creating a cachable list of program identifiers, containing a copy of said designated program image and each of said library programs identified in said fourth list; and

returning said cachable list of program identifiers in response to said link method invocation.

30. The method of claim 29 wherein said programs are objects and said computing system is an object oriented system.

31. The method of claim 30 wherein said identifiers and said program identifiers are memory objects.

32. A Dynamic Linking system for a data processing system, said data processing system having application programs, library programs and an operating system, said Dynamic Linking system comprising:

an image cache for caching images of linked application programs and said application program's related library programs;

a first searching facility, coupled to said image cache, for searching said image cache to determine whether said image cache contains a linked application program image which matches a designated application program, said search being made in response to a request to link said designated application program; and

a first communications facility, coupled to said first searching facility, for responding to said request to link said designated application program, said response containing a list of program identifiers and addresses related to said program identifiers obtained from said image cache in the event that a linked application program image is found in said image cache which matches said designated application program.

33. The Dynamic Linking system as recited in claim 32 further comprising:
a library program cache, coupled to said image cache, for caching library programs; and
a second searching facility, coupled to said library program cache, for searching said library program cache to determine whether said library program cache contains a library program which matches a designated library program which is used as a search argument.
34. The Dynamic Linking system as recited in claim 33 wherein said search of said library cache is made in the event that a linked application program image is not found in said image cache which matches said designated application program, and wherein said first searching facility obtains a copy of an unlinked application program image which matches said designated application program.
35. The Dynamic Linking system as recited in claim 34 further comprising:
a first linking facility, coupled to said second searching facility, for linking a found library program to said unlinked application program image, in the event that said found library program is found in said library program cache which matches said designated library program.
36. The Dynamic Linking system as recited in claim 35 further comprising:
a third searching facility, coupled to said second searching facility, for finding unlinked library programs in the event that a library program is not found in said library program cache which matches said designated library program; and
a second linking facility, coupled to said third searching facility, for linking said unlinked library program to each other and to said found library programs and to said unlinked application program, producing a new linked application program image with related library programs.
37. The Dynamic Linking system as recited in claim 36 wherein:
said second linking facility partially links any of said unlinked library programs and stores said partially linked library programs in said library program cache; and
said second linking facility stores said new linked application program image with related library programs in said image cache.
38. The Dynamic Linking system as recited in claim 37 wherein said library programs in said library program cache are partially linked library programs.
39. The Dynamic Linking system as recited in claim 38 wherein said images of linked application programs and said application program's related library programs in said image cache are fully linked application programs and said application program's related library programs.
40. The Dynamic Linking system as recited in claim 39 wherein said application programs and library programs are object oriented programs.
41. A Dynamic Linking system for a data processing system, said data processing system having application programs, library programs and an operating system, said Dynamic Linking system comprising:
a library program cache for caching partially linked library programs;
a first searching facility, coupled to said library program cache, for searching said library program cache to determine whether said library program cache contains a partially linked library program which is required by a designated application program, said search being made in response to a request to link said designated application program; and
a first linking facility, coupled to said first searching facility, for linking a found library program to said designated application program, in the event that said found library program is found in said library program cache.
42. The Dynamic Linking system as recited in claim 41 wherein said first searching facility further comprises:
a second searching facility, coupled to said library program cache, for searching said library program cache to determine whether said library program cache contains a library program which matches a designated library program.
43. The Dynamic Linking system as recited in claim 42 further comprising:
a third searching facility, coupled to said second searching facility, for finding unlinked library programs in the event that a library program is not found in said library program cache which matches said designated library program; and
a second linking facility, coupled to said

third searching facility, for linking said unlinked library programs to each other and to said found library programs and to said designated application program, producing a new linked application program image with related library programs. 5

44. The Dynamic Linking system as recited in claim 43 wherein:
said second linking facility partially links 10
any of said unlinked library programs and stores said partially linked library programs in said library program cache.
45. The Dynamic Linking system as recited in 15
claim 44 wherein said library programs in said library program cache are partially linked library programs.
46. The Dynamic Linking system as recited in 20
claim 45 wherein said application programs and said library programs are object oriented programs.

25

30

35

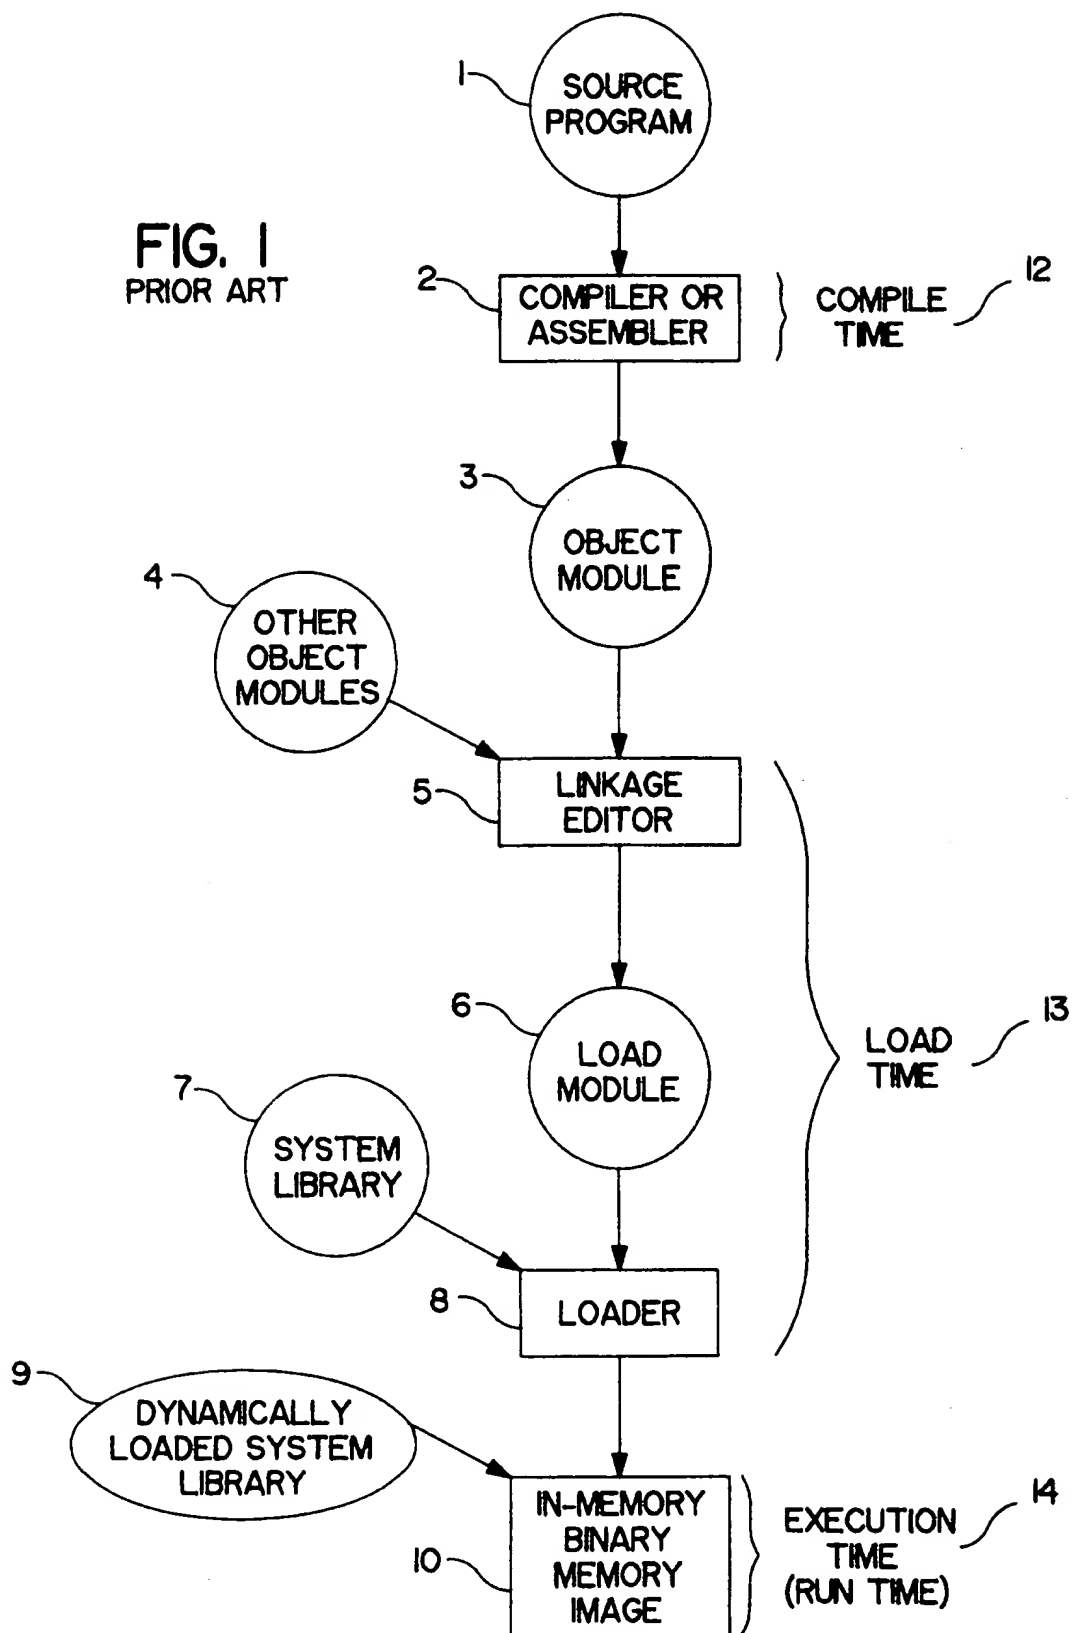
40

45

50

55

FIG. 1
PRIOR ART






FIG. 2a
PRIOR ART


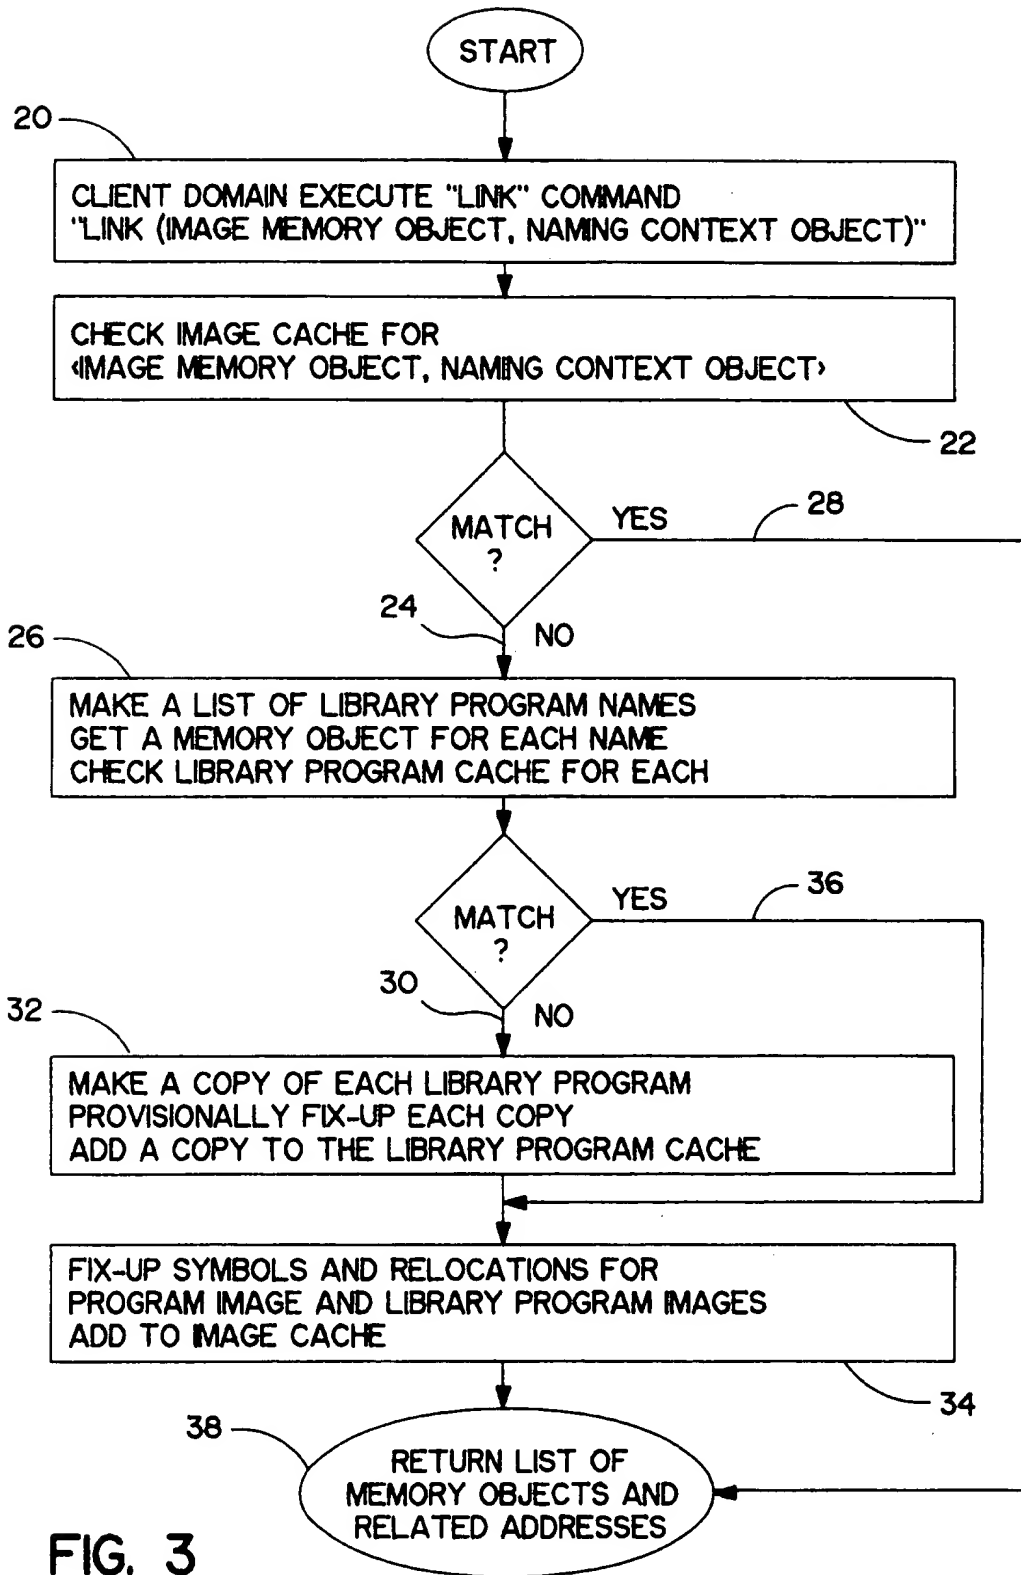


FIG. 2b



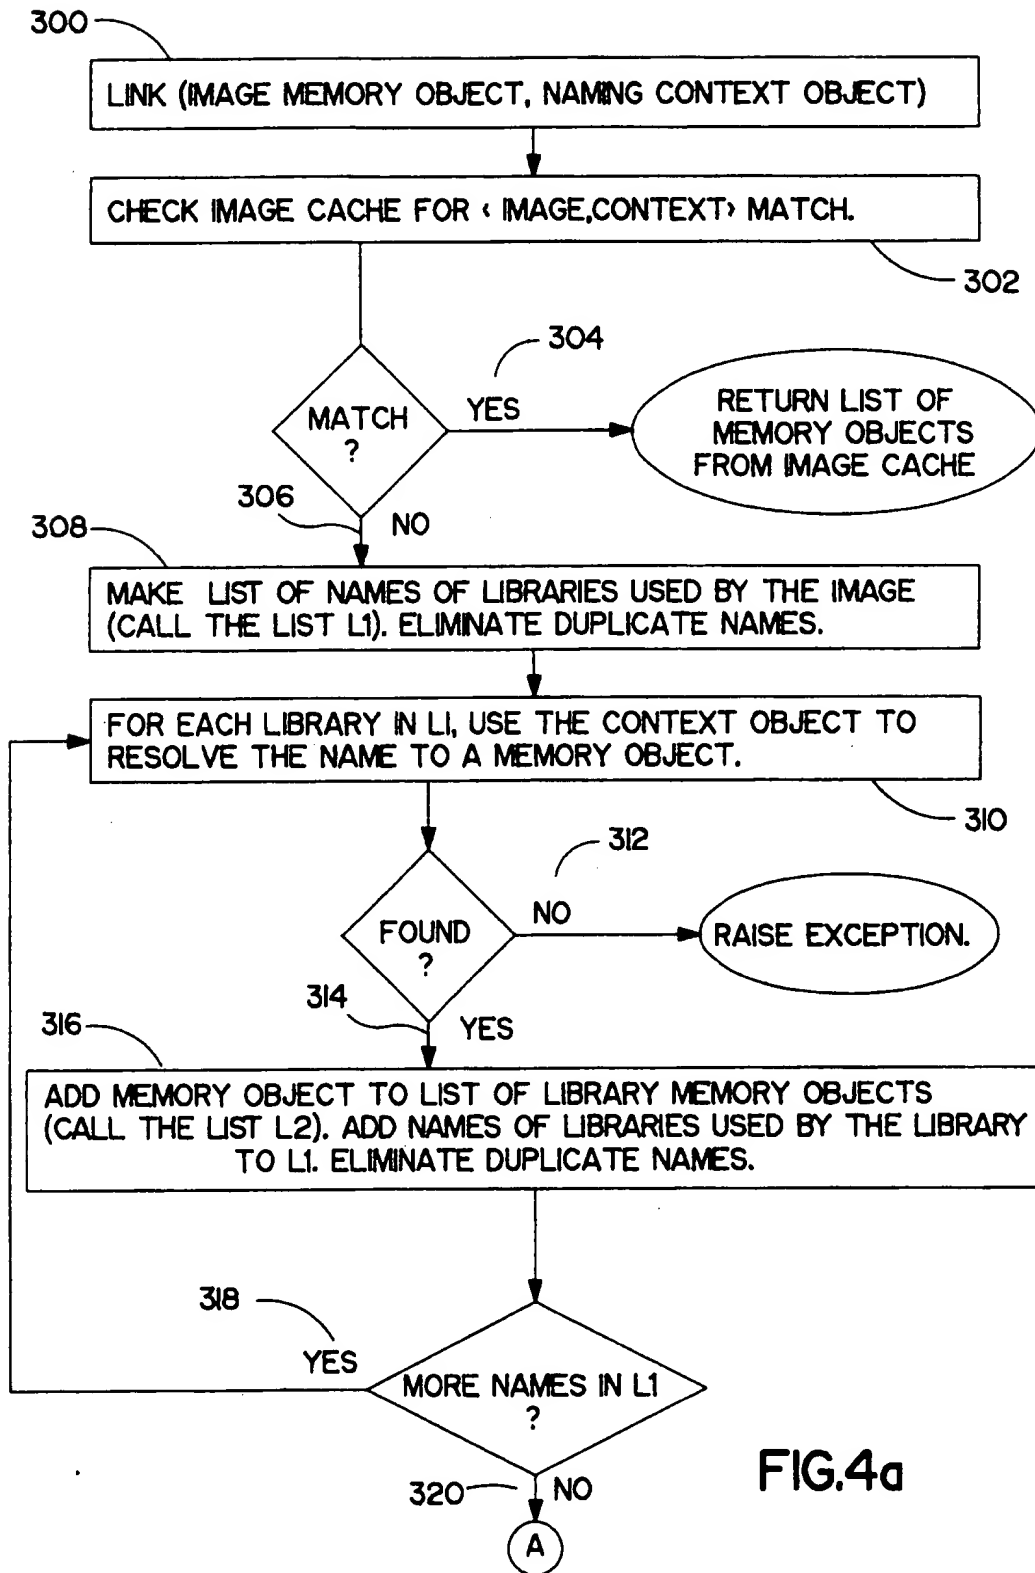


FIG. 4a

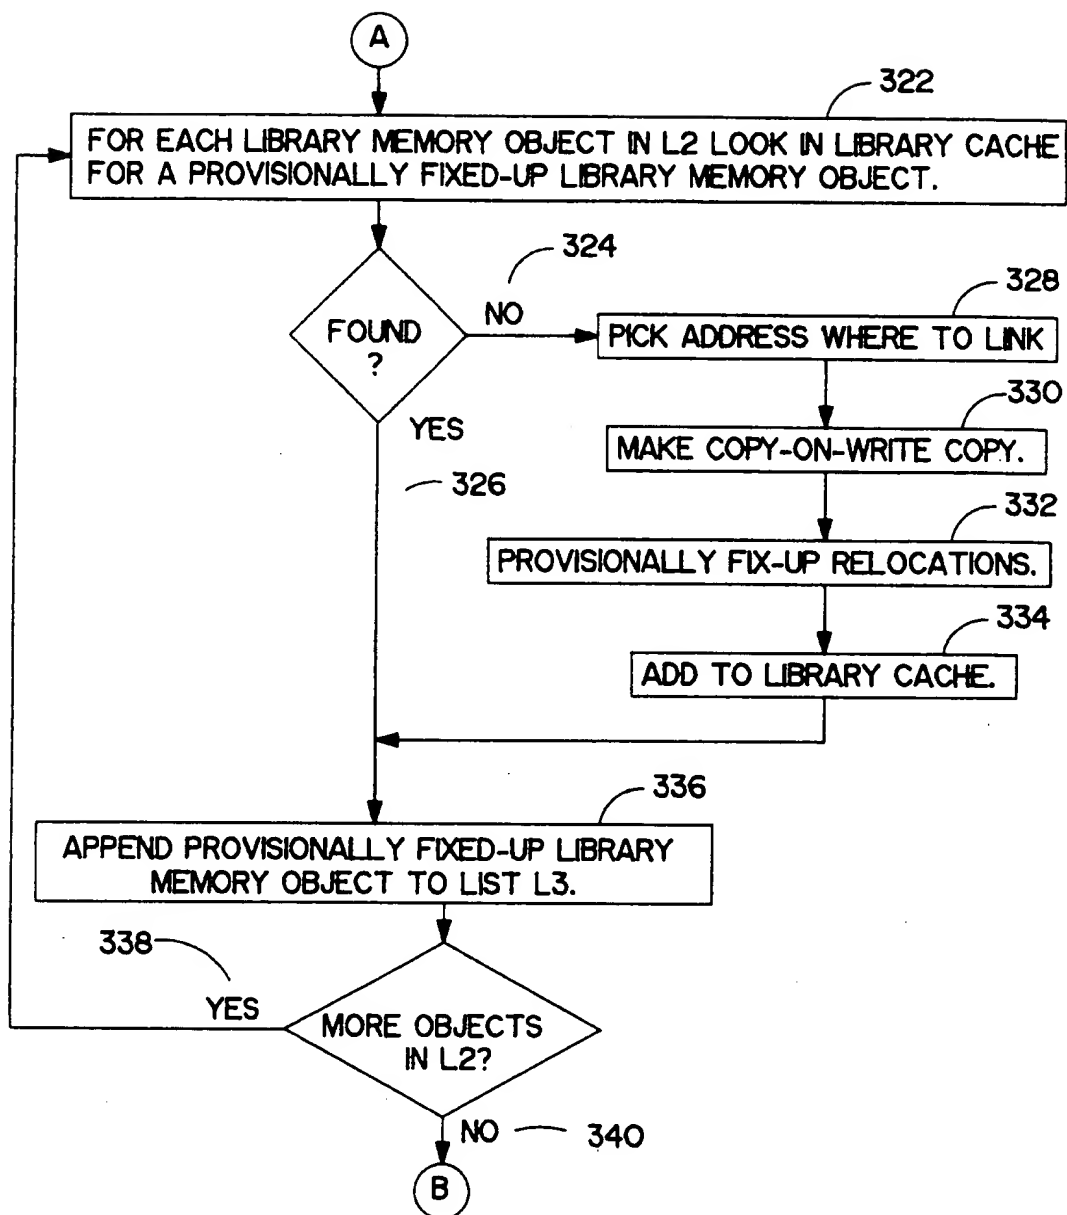


FIG. 4b

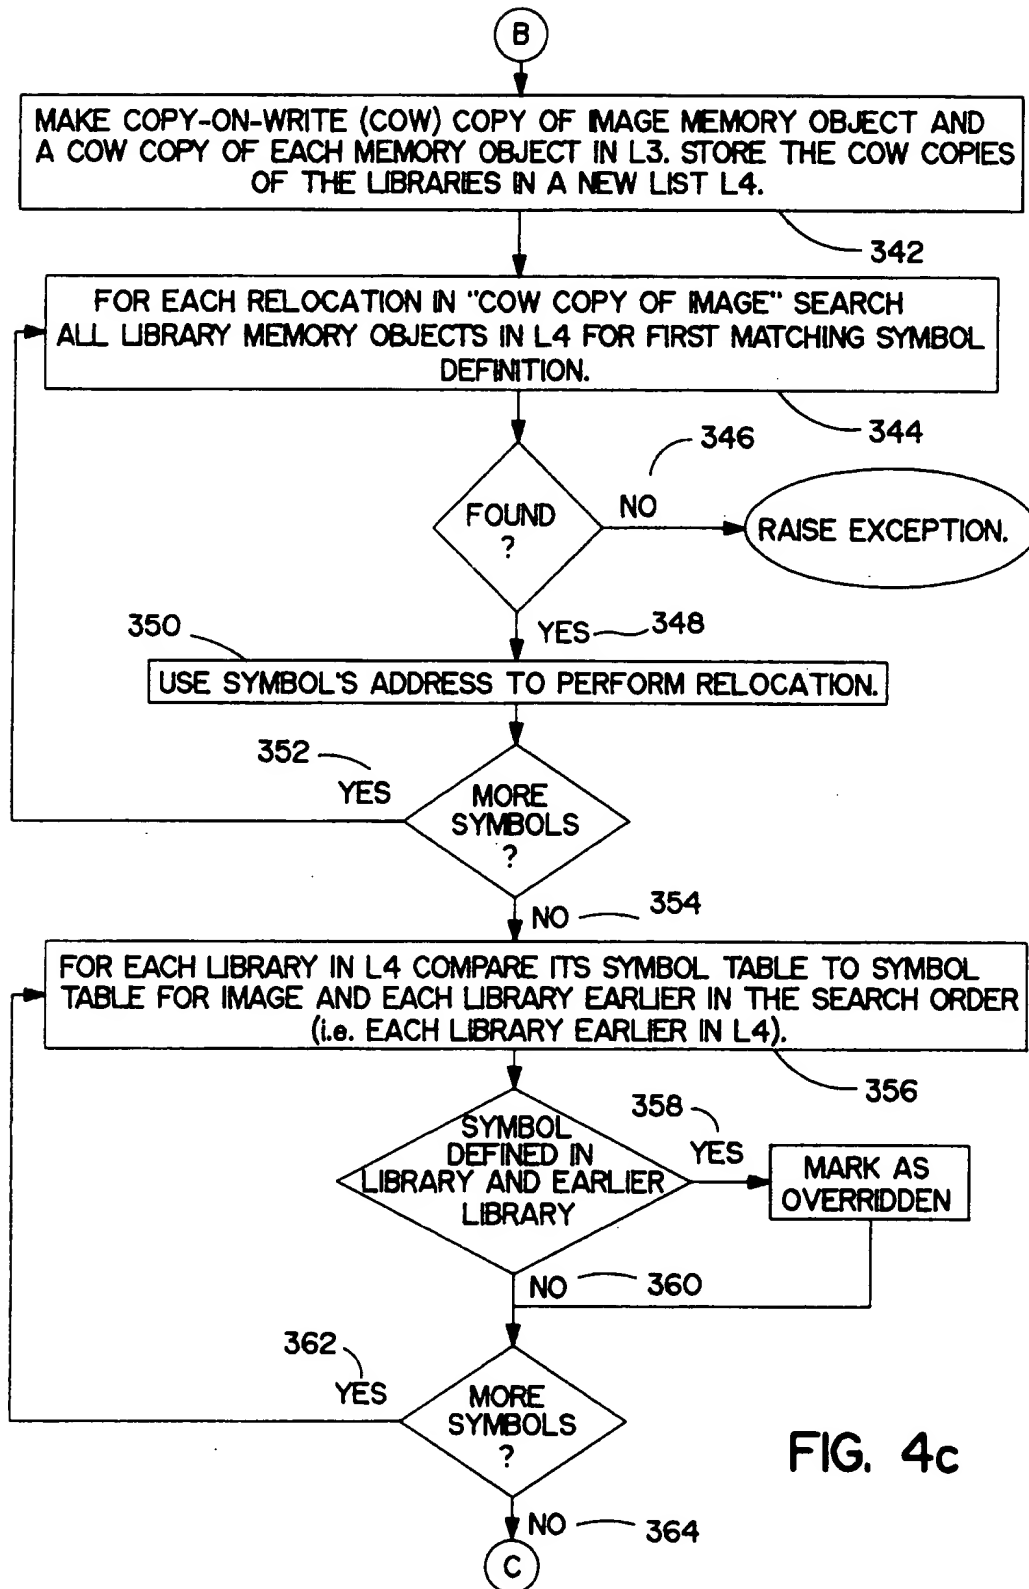


FIG. 4c

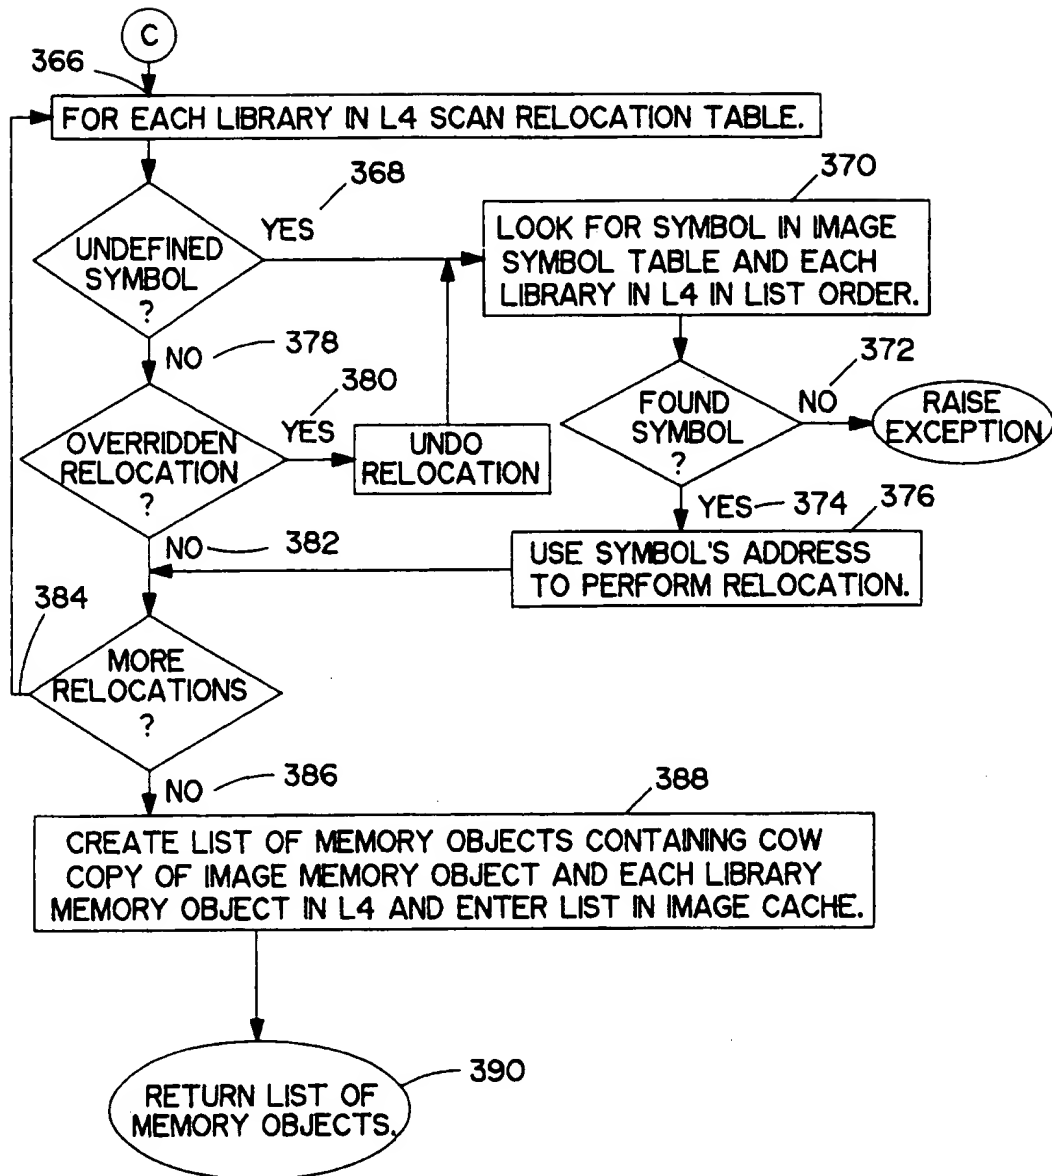


FIG. 4d




FIG. 5

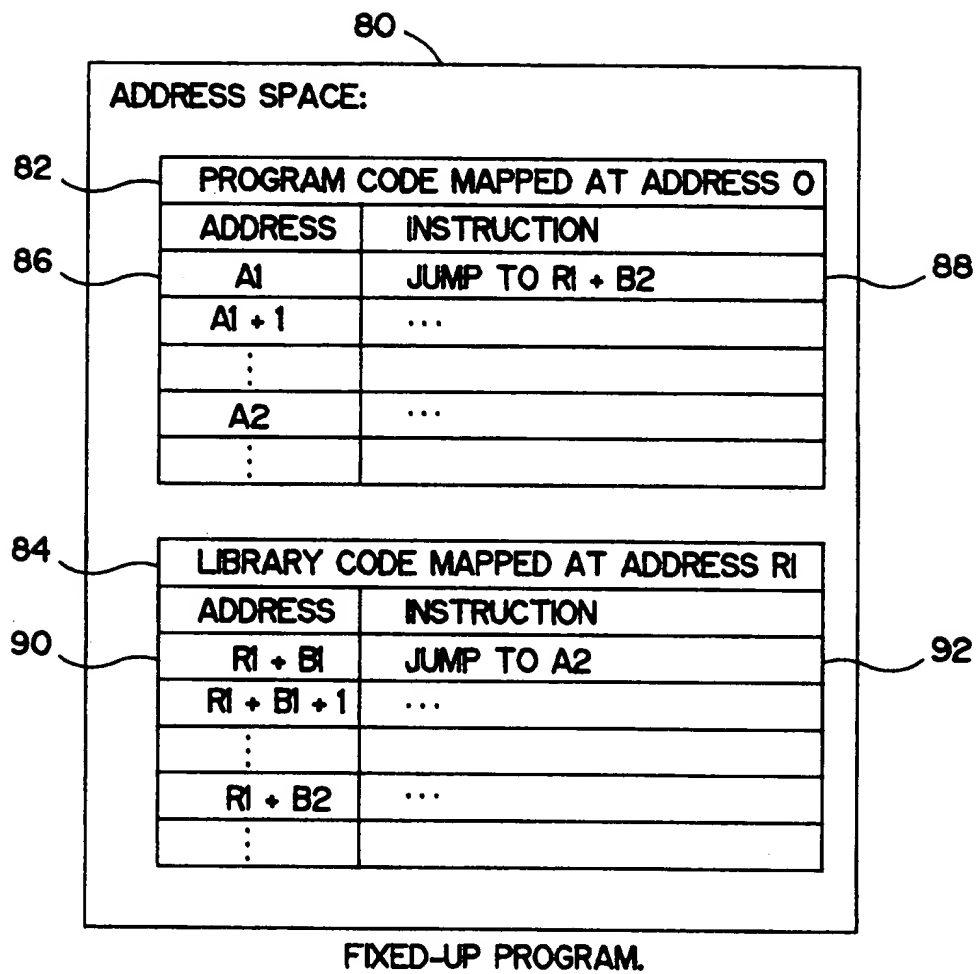
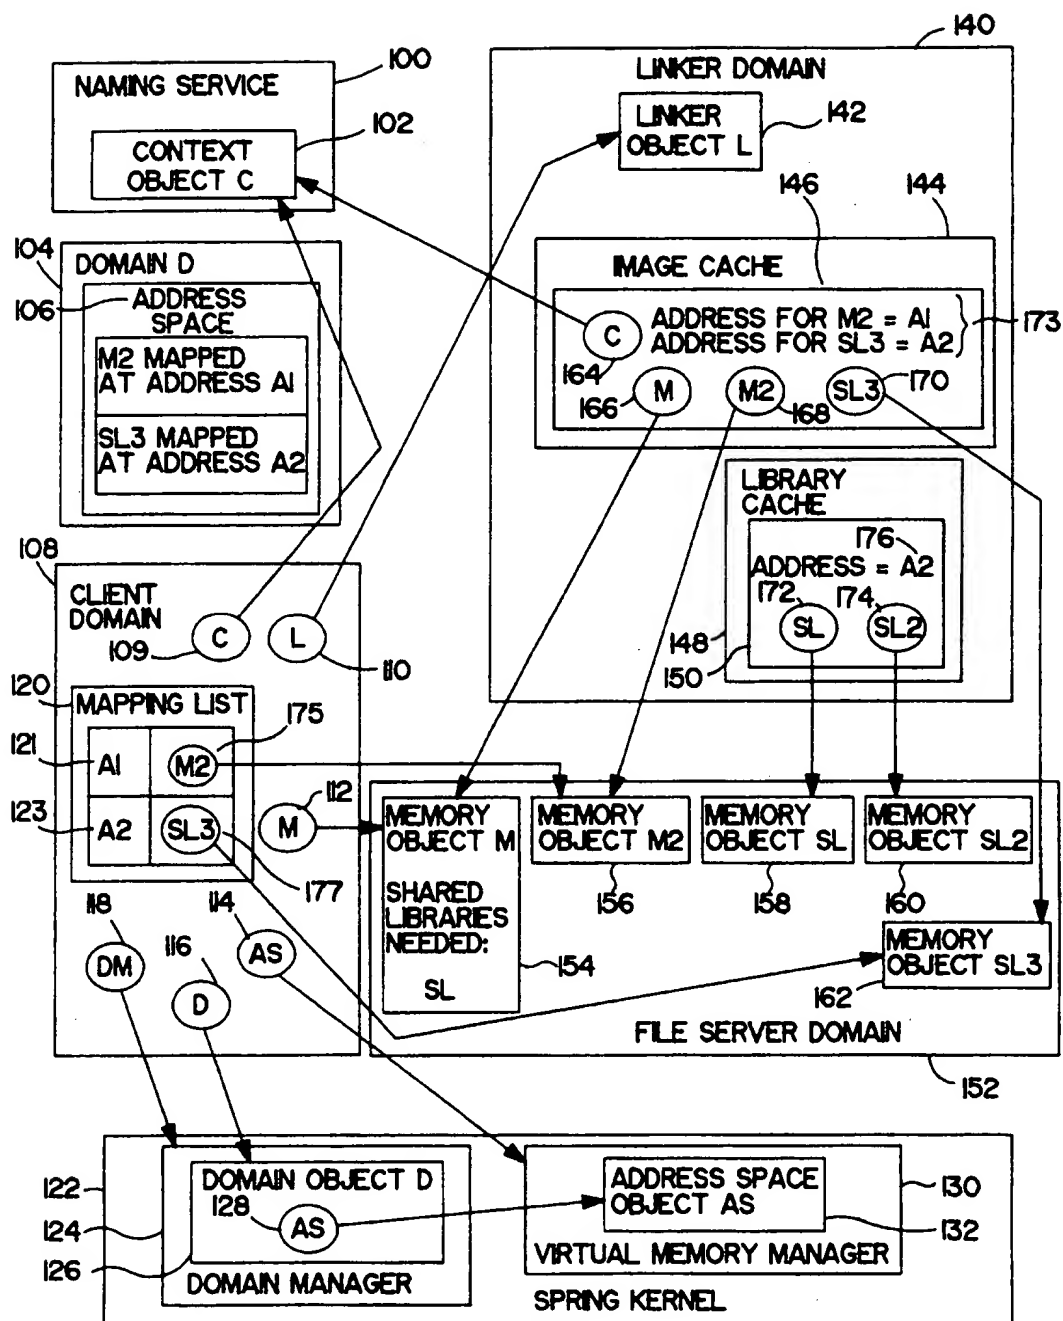


FIG. 6



EXAMPLE OF LINKING IN SPRING

FIG. 7



EUROPEAN PATENT APPLICATION

(51) Int. Cl.⁶: G06F 9/445

(43) Date of publication A2:
19.10.1994 Bulletin 1994/42

(21) Application number: 94301785.5

(22) Date of filing: 14.03.1994

(84) Designated Contracting States:
DE FR IT NL SE

(30) Priority: 13.04.1993 US 46827

(71) Applicant: **SUN MICROSYSTEMS, INC.**
Mountain View, California 94043-1100 (US)

(72) Inventors:

- Nelson, Michael N.
San Carlos, California 94070 (US)
- Hamilton, Graham
Palo Alto, California 94303 (US)

**(74) Representative: Wombwell, Francis
Potts, Kerr & Co.
15, Hamilton Square
Birkenhead Merseyside L41 6BR (GB)**

(54) High performance dynamic linking through caching

(57) A method is disclosed for a method to dynamically link a new program image and related library programs into an executable application program image. The method provides for producing a linked list of the required programs by giving the linker a n argument representing the designated program image and a naming context which contains data on the associated library programs which are to be linked together. The linker finds all of the required programs, and links them together. The parent maps the program images into the designated addresses thereby completing the linking of the executable application program. In finding the required programs, the linker first checks the image cache to see if the new program and its related library programs is already linked and cached because it was executed before. If the new program is not found in the image cache, the linker object then checks the library program cache to see if the library programs are cached in partially linked form, and will use as many of these as it can find. For any library programs that must still be located, the linker will retrieve them from a data store and will then proceed to link all library programs and the new program image together to form an executable whole. This process of caching new program images with their library programs and caching partially linked library programs individually guarantees a procedure which minimizes the time delay in program start-up when a new program is executed.

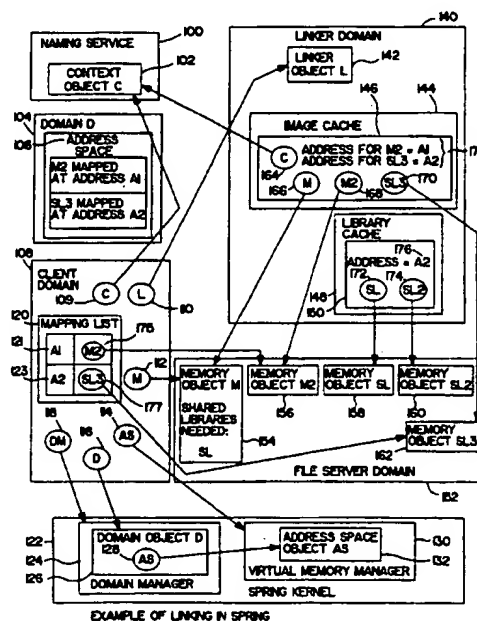


FIG. 7



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 94 30 1785

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cls)
A	SOFTWARE PRACTICE & EXPERIENCE, vol. 21, no. 4, 1 April 1991 pages 375-390, XP 000147180 WILSON HO W ET AL 'AN APPROACH TO GENUINE DYNAMIC LINKING' * page 384, line 1 - line 13 *	1-46	G06F9/445
A	IBM JOURNAL OF RESEARCH AND DEVELOPMENT, vol. 34, no. 1, 1 January 1990 pages 98-103, XP 000128185 AUSLANDER M A 'MANAGING PROGRAMS AND LIBRARIES IN AIX VERSION 3 FOR RISC SYSTEM/6000 PROCESSORS' * page 99, right column, line 3 - line 39 *	1-46	
A	PATENT ABSTRACTS OF JAPAN vol. 016 no. 495 (P-1436) ,14 October 1992 & JP-A-04 178731 (HITACHI LTD) 25 June 1992, * abstract *	1-46	
The present search report has been drawn up for all claims			TECHNICAL FIELDS SEARCHED (Int.Cls)
			G06F
Place of search		Date of completion of the search	Examiner
THE HAGUE		29 February 1996	Brandt, J
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons</p> <p>Δ : member of the same patent family, corresponding document</p>			

EPO FORM 1503 01.02 (Pct/OI)